

Parametric-Rate Inference for One-Sided Differentiable Parameters

Alexander R. Luedtke*

and

Mark J. van der Laan†

Division of Biostatistics, University of California, Berkeley

May 26, 2016

Abstract

Suppose one has a collection of parameters indexed by a (possibly infinite dimensional) set. Given data generated from some distribution, the objective is to estimate the maximal parameter in this collection evaluated at this distribution. This estimation problem is typically non-regular when the maximizing parameter is non-unique, and as a result standard asymptotic techniques generally fail in this case. We present a technique for developing parametric-rate confidence intervals for the quantity of interest in these non-regular settings. We show that our estimator is asymptotically efficient when the maximizing parameter is unique so that regular estimation is possible. We apply our technique to a recent example from the literature in which one wishes to report the maximal absolute correlation between a prespecified outcome and one of p predictors. The simplicity of our technique enables an analysis of the previously open case where p grows with sample size. Specifically, we only require that $\log p$ grows slower than \sqrt{n} , where n is the sample size. We show that, unlike earlier approaches, our method scales to massive data sets: the point estimate and confidence intervals can be constructed in $O(np)$ time.

Keywords: stabilized one-step estimator; non-regular inference; variable screening.

*This author gratefully acknowledges the support of the Berkeley Fellowship.

†This author gratefully acknowledges the support of NIH grant R01 AI074345-06.

1 Introduction

Many semiparametric and nonparametric estimation problems yield estimators which achieve a parametric rate of convergence. These estimators are often asymptotically linear, in that they can be written as an empirical mean of an influence function applied to the data. Valid choices of the influence function can be derived as gradients for a functional derivative of the parameter of interest. Applying the central limit theorem then immediately yields Wald-type confidence intervals which achieve the desired parametric rate. Such problems have been studied in depth over the past several decades [Pfanzagl, 1990, van der Vaart, 1991, Bickel et al., 1993, van der Laan and Robins, 2003].

While remarkably general, these approaches rely on the key condition that the parameter of interest is sufficiently differentiable for such a gradient to exist. Statisticians are increasingly encountering problems for which parametric-rate estimation is theoretically possible but the parameter is insufficiently differentiable to yield a standard first-order expansion demanded by older techniques. For example, suppose we observe baseline covariates, a binary treatment, and an outcome occurring after treatment. We wish to learn the mean outcome under the optimal individualized treatment strategy, i.e. the treatment strategy which makes treatment decisions which are allowed to use baseline covariate information to make treatment decisions Chakraborty and Moodie [2013]. As another example, suppose we observe a vector of covariates (X_1, \dots, X_p) and an outcome Y . We wish give a confidence interval the maximal absolute correlation between a covariate X_k and Y , or at least a lower bound on this quantity since this will suffice for a variable screening procedure. Alternatively, we may only wish to test the null hypothesis that the maximal absolute correlation is zero. McKeeague and Qian [2015] provide a test of this null hypothesis using an adaptive resampling test (ART), a framework initially introduced in Laber and Murphy [2011] for estimating classification error.

These problems belong to a larger class of problems in which one observes O_1, \dots, O_n drawn independently from a P_0 in some (possibly nonparametric) statistical model \mathcal{M} and wishes to estimate

$$\Psi_n(P) \equiv \max_{d \in \mathcal{D}_n} \Psi^d(P), \tag{1}$$

at $P = P_0$, where \mathcal{D}_n is an index set that may rely on sample size and each Ψ^d is a sufficiently

differentiable parameter to permit parametric-rate estimation using classical methods such as those presented in Bickel et al. [1993]. When there is no unique maximizer $d \in \mathcal{D}_n$ of $\Psi^d(P)$, then the inference problem is typically non-regular, in the sense that the parameter $P \mapsto \max_{d \in \mathcal{D}_n} \Psi^d(P)$ is not sufficiently differentiable to allow the use of standard techniques for obtaining root- n inference. In these cases, the parameter of interest is termed one-sided pathwise differentiable [Hirano and Porter, 2012]. In univariate calculus, functions such as $f(x) = \max\{x, 0\}$ are one-sided differentiable at zero in that the left and right limits of $[f(x + \epsilon) - f(x)]/\epsilon$ are well-defined but disagree. The same holds for the Ψ_n evaluated at a distribution P_0 , but now the one-sided differentiability is caused by the subset of \mathcal{D}_n containing the indices which maximize the expression on the right in (1). A small fluctuation in P_0 can greatly reduce the subset of maximizing indices, leading to different derivatives depending on the fluctuation taken.

In this work, we present a method which, loosely, splits the sample in such a way that the estimated index in \mathcal{D}_n which maximizes $\Psi^d(P_0)$ is conditioned on so that this estimated index need not have a limit. We do this iteratively to ensure that our estimator gets the full benefit of the sample size n . When the parameter is fixed with sample size and the d maximizing $\Psi^d(P_0)$ is fixed, we show that our estimator is asymptotically efficient, and therefore also regular. Thus our estimator adapts to the non-regularity of the estimation problem.

Our estimator is inspired by the online estimator for pathwise differentiable parameters presented in van der Laan and Lendle [2014] and a subsequent modification of this estimator in Luedtke and van der Laan [2016] to deal with the non-regularity when estimating the mean outcome under an optimal treatment rule. Such estimators are designed to be efficient in both computational complexity and storage requirements. We show that the estimator that we present in this work inherits many of these computational efficiency properties. We apply our technique to estimate the maximal absolute correlation considered in McKeague and Qian [2015]. In this problem, we show that our estimator runs efficiently in both dimension and sample size, with a runtime of $O(np)$. In practice, this means that the lead author can implement our estimator using only R code and screen $p = 100\,000$ variables using $n = 1\,000$ samples on a single core of his laptop in under a minute. Thus our estimator seems to have both the statistical efficiency that has been demanded of estimators for generations and

the computational efficiency that is becoming increasingly important in this new big data era.

2 Estimator

We will now present our technique for a general estimation problem. Before doing so, we must introduce the notion of pathwise differentiability, since this provides the key object needed to construct our estimator.

2.1 Pathwise differentiability

We assume that each parameter Ψ^d , $d \in \mathcal{D}_n$ for any n , is pathwise differentiable for all distributions in our model [see, e.g., Pfanzagl, 1990, Bickel et al., 1993]. For each $P \in \mathcal{M}$, we let $D^d(P)$ denote the canonical gradient of Ψ^d at P . By definition $D^d(P)(O)$ is mean zero with finite variance under sampling from P . Typically pathwise differentiability implies that Ψ^d satisfies the following linear expansion for any $P \in \mathcal{M}$ and $d \in \mathcal{D}_n$:

$$\Psi^d(P) - \Psi^d(P_0) = - \int D^d(P)(o) dP_0(o) + \text{Rem}_n^d(P), \quad (2)$$

where we omit the dependence of $\text{Rem}_n^d(P)$ on P_0 in the notation and indicate its possible dependence on sample size with the subscript n . Above $\text{Rem}_n^d(P)$ is a second-order remainder term that is small whenever P is close to P_0 . We consider this condition more closely in our examples, but for non-sample size dependent parameters this term can typically be made to be $O_{P_0}(1/n)$ in a parametric model and often can be made to be $o_{P_0}(1/\sqrt{n})$ in a nonparametric model. For a more thorough presentation, see Pfanzagl [1990] or Bickel et al. [1993].

2.2 Estimator and confidence interval

We now present a stabilized one-step estimator for problems of the type found in (1) when the required differentiability condition on Ψ^d holds.

Let $\{\ell_n\}$ be some sequence such that $n - \ell_n \rightarrow \infty$. One possible choice is $\ell_n = 0$ for all n . For each $j = \ell_n, \dots, n - 1$, let d_{nj} represent an estimate of a maximizer of (1), \hat{P}_{nj} be

an estimate of the likelihood obtained using observations $(O_i : i = 1, \dots, j)$, and \hat{D}_{nj} equal $D^d(P)$ evaluated at $P = \hat{P}_{nj}$ and $d = d_{nj}$. For nonnegative weights $w_{\ell_n}, \dots, w_{n-1}$ that we will define shortly with $\sum_{j=\ell_n}^{n-1} w_j = n - \ell_n$, our stabilized one-step estimator takes the form

$$\psi_n \equiv \frac{1}{n - \ell_n} \sum_{j=\ell_n}^{n-1} w_j \left[\Psi^{d_{nj}}(\hat{P}_{nj}) + \hat{D}_{nj}(O_{j+1}) \right].$$

Our proposed 95% confidence interval has the form

$$[\text{LB}_n, \text{UB}_n] \equiv \left[\psi_n \pm 1.96 \frac{\bar{\sigma}_n}{\sqrt{n - \ell_n}} \right],$$

where we will define $\bar{\sigma}_n$ momentarily and one can replace 1.96 by the desired quantile of the normal distribution to modify the confidence level.

We now define the weights. Let $\hat{\sigma}_{nj}^2$ represent an estimate of the variance of $\hat{D}_{nj}(O)$, $O \sim P_0$, conditional on observations O_1, \dots, O_j . This estimate should only rely on those j observations. Often we can let

$$\hat{\sigma}_{nj}^2 \equiv \frac{1}{j} \sum_{i=1}^j \left[\hat{D}_{nj}(O_i) - \frac{1}{j} \sum_{i=1}^j \hat{D}_{nj}(O_i) \right]^2, \quad j = \ell_n, \dots, n-1.$$

The standard deviation type variable in the confidence interval definition is given by $\bar{\sigma}_n \equiv \left(\frac{1}{n - \ell_n} \sum_{j=\ell_n}^{n-1} \hat{\sigma}_{nj}^{-1} \right)^{-1}$, and the weights are given by $w_j \equiv \bar{\sigma}_n \hat{\sigma}_{nj}^{-1}$, where we have omitted the possible dependence of the weights on sample size in the notation.

Our estimator ψ_n is similar to the online one-step estimator developed in van der Laan and Lendle [2014] for streaming data, but it weights each term proportionally to the estimated inverse standard deviation of $\hat{D}_{nj}(O)$ when $O \sim P_0$. Our confidence interval takes a form similar to a Wald-type confidence interval, but replaces the typical standard deviation with $\bar{\sigma}_n$ and has width on the order of $1/\sqrt{n - \ell_n}$ rather than $1/\sqrt{n}$. Note of course that $\ell_n = o(n)$ implies that $1/\sqrt{n - \ell_n} - 1/\sqrt{n}$ converges to zero.

2.3 First main result: validity of confidence interval

We now prove the validity of our confidence interval. Let $\sigma_{nj}^2 \equiv \text{Var}_{P_0} \left(\hat{D}_{nj}(O) | O_1, \dots, O_j \right)$. The validity of the lower bound of the confidence interval relies on the following conditions:

- C1) There exists some $M < \infty$ such that $\frac{1}{n - \ell_n} \sum_{j=\ell_n}^{n-1} P_0 \left(\frac{|\hat{D}_{nj}(O)|}{\hat{\sigma}_{nj}} < M \mid O_0, \dots, O_{j-1} \right) \rightarrow 1$ in probability as $n \rightarrow \infty$.

- C2) $\frac{1}{n-\ell_n} \sum_{j=\ell_n}^{n-1} \left| \frac{\sigma_{nj}^2}{\hat{\sigma}_{nj}^2} - 1 \right| \rightarrow 0$ in probability as $n \rightarrow \infty$.
C3) $\frac{1}{\sqrt{n-\ell_n}} \sum_{j=\ell_n}^{n-1} \hat{\sigma}_{nj}^{-1} \widehat{\text{Rem}}_{nj} \rightarrow 0$ in probability as $n \rightarrow \infty$, where $\widehat{\text{Rem}}_{nj} \equiv \text{Rem}^{d_{nj}}(\hat{P}_{nj})$.

The validity of the upper bound requires the following additional condition:

- C4) $\frac{1}{\sqrt{n-\ell_n}} \sum_{j=\ell_n}^{n-1} \hat{\sigma}_{nj}^{-1} [\Psi^{d_{nj}}(P_0) - \Psi_n(P_0)]$ converges to zero in probability as $j \rightarrow \infty$.

We now present our main result. We discuss the conditions in Section 2.5.

Theorem 1 (Validity of confidence interval). *If C1), C2), and C3) hold, then*

$$\liminf_{n \rightarrow \infty} \Pr(\Psi_n(P_0) \geq \text{LB}_n) \geq 1 - \alpha/2.$$

If C4) also holds, then

$$\lim_{n \rightarrow \infty} \Pr(\text{LB}_n \leq \Psi_n(P_0) \leq \text{UB}_n) = 1 - \alpha.$$

Proof. The definition ψ_n combined with (2) yield that

$$\begin{aligned} \sqrt{n-\ell_n} \bar{\sigma}_n^{-1} [\psi_n - \Psi_n(P_0)] &= \frac{1}{\sqrt{n-\ell_n}} \sum_{j=\ell_n}^{n-1} \hat{\sigma}_{nj}^{-1} \left(\hat{D}_{nj}(O_{j+1}) - E_{P_0} [\hat{D}_{nj}(O) | O_1, \dots, O_j] \right) \\ &+ \frac{1}{\sqrt{n-\ell_n}} \sum_{j=\ell_n}^{n-1} \hat{\sigma}_{nj}^{-1} [\Psi^{d_{nj}}(P_0) - \Psi_n(P_0) + \widehat{\text{Rem}}_{nj}]. \end{aligned} \quad (3)$$

The second line converges to zero in probability by C3) and C4). By C1), C2), and the martingale central limit theorem for triangular arrays in Gaenssler et al. [1978], (3) converges in distribution to a standard normal random variable. A standard Wald-type confidence interval construction argument shows that the confidence interval has coverage approaching $1 - \alpha$ under C1) through C4).

Now suppose C4) does not hold. By (1), $\sum_{j=\ell_n}^{n-1} \hat{\sigma}_{nj}^{-1} [\Psi^{d_{nj}}(P_0) - \Psi_n(P_0)] \leq 0$. The same argument readily shows the validity of the lower bound under only C1), C2), and C3). \square

2.4 Second main result: efficiency when the maximizer in (1) is unique

We have presented a parametric-rate estimator for $\Psi_n(P_0)$, but thus far we have not made any claims about the efficiency of our estimator. In this section, we consider a fixed parameter in

(1) that does not rely on sample size. We therefore omit the n subscript in many quantities to indicate their lack of dependence on sample size. We will give conditions under which our estimator is asymptotically efficient among all regular, asymptotically linear estimators. The efficiency bound is not typically well-defined when the maximizer is non-unique due to the non-regularity of the problem – generally in this case no regular, asymptotically linear estimator exists, so neither does an efficient member of this class [Hirano and Porter, 2012]. Thus the conditions that we give in this section will typically only hold when the maximizer $d_0 \in \mathcal{D}$ in (1) is unique.

We use the following additional assumptions for our efficiency result:

$$\text{C5)} \quad E_{P_0} \left[\left(\hat{D}_j(O) - D^{d_0}(P_0)(O) \right)^2 \middle| O_1, \dots, O_j \right] \rightarrow 0 \text{ in probability as } j \rightarrow \infty.$$

C6) There exists some $M < \infty$ such that $P_0(D^{d_0}(P_0)(O) < M)$ and $P_0(\hat{D}_j(O) < M)$ with probability approaching 1 as $j \rightarrow \infty$.

C7) $\inf_{j \geq 1} \hat{\sigma}_j^2 > \gamma$ with probability 1 over draws of $(O_j : j = 0, 1, \dots)$.

We discuss the conditions immediately following the theorem.

Theorem 2 (Asymptotic efficiency). *Suppose that Ψ does not depend on sample size and is pathwise differentiable with canonical gradient $D^{d_0}(P_0)$. Further suppose that $\ell_n = o(n)$. If C1) through C7) hold, then*

$$\bar{\sigma}_n^2 \rightarrow \text{Var}_{P_0}(D^{d_0}(P_0)(O)) \text{ in probability as } n \rightarrow \infty.$$

Furthermore,

$$\psi_n - \Psi(P_0) = \frac{1}{n} \sum_{i=1}^n D^{d_0}(P_0)(O_i) + o_{P_0}(n^{-1/2}).$$

Thus ψ_n is asymptotically efficient among all regular, asymptotically linear estimators.

The proof is entirely analogous to the proof of Corollary 3 in Luedtke and van der Laan [2016] so is omitted.

The additional conditions needed for this result over Theorem 1 are mild when the maximizing index is unique. Condition C5) says that Ψ should have the same canonical gradient as Ψ^{d_0} . While this should be manually checked in each example, it will be fairly typical when the maximizer is unique, since in this case an arbitrarily small fluctuation of P_0 will

generally not change the maximizer. This is similar to problems in introductory calculus where the derivative at the maximum is zero. Condition C5) requires that $\hat{D}^j(O)$ converge to $D^{d_0}(P_0)(O)$ in mean-squared error, which is to be expected if \hat{P}_{nj} begins to approximate P_0 and d_{nj} converges to the unique maximizer d_0 as $n, j \rightarrow \infty$. Condition C6) is a bounding assumption on the canonical gradient and estimates thereof that will hold in many examples of interest. Finally, Condition C7) will hold if one knows that $\text{Var}_{P_0} [D^d(P)(O)]$ is bounded away from zero uniformly in $P \in \mathcal{M}$ and $d \in \mathcal{D}$, and uses this knowledge to truncate $\hat{\sigma}_j^2$ at $\gamma_j > 0$ for some deterministic sequence $\gamma_j \rightarrow 0$. For γ_j sufficiently small and j sufficiently large this truncation scheme will then have no effect on the variance estimates $\hat{\sigma}_j^2$.

2.5 Discussion of conditions of Theorem 1

In this section, we again consider the setting where the parameter does not depend on sample size, and consequently omit the n subscript to quantities which no longer depend on sample size. We will show that C7) and the following conditions imply the conditions of Theorem 1:

C9) $\hat{\sigma}_j^2 - \sigma_j^2$ converges to zero in probability as $j \rightarrow \infty$.

C10) $\sqrt{j} \widehat{\text{Rem}}_j \equiv \sqrt{j} \text{Rem}^{d_j}(\hat{P}_j)$ converges to zero in probability as $j \rightarrow \infty$.

The validity of the upper bound requires the following additional condition:

C11) $\sqrt{j} [\Psi^{d_j}(P_0) - \Psi(P_0)]$ converges to zero in probability as $j \rightarrow \infty$.

For simplicity, we will take $\ell_n = 0$ in this section.

We now discuss the conditions. Condition C1) is an immediate consequence of C7) and $D^d(P)(o)$ being uniformly bounded in $P \in \mathcal{M}$, $d \in \mathcal{D}$, $o \in \mathcal{O}$. This will be plausible in many situations, including the examples in this paper. A more general Lindeberg-type condition also suffices [see Condition C1 in Luedtke and van der Laan, 2016], though we omit its presentation here for brevity.

The other three conditions all rely on terms like $\frac{1}{n} \sum_{j=0}^{n-1} R_j$ converging to zero in probability, possibly at some rate. Ideally we want a stochastic version of the fact that, for $\beta \in [0, 1)$,

$$\frac{1}{n} \sum_{j=1}^n j^{-\beta} \approx \frac{1}{n} \int_1^n j^{-\beta} dj \approx \frac{n^{-\beta}}{1-\beta} \text{ when } n \text{ is large.} \quad (4)$$

Lemma 6 of Luedtke and van der Laan [2016] establishes this result. We restate it here for convenience.

Lemma 3 (Lemma 6 in Luedtke and van der Laan, 2016). *Suppose that R_j is some sequence of (finite) real-valued random variables such that $R_j = o_{P_0}(j^{-\beta})$ for some $\beta \in [0, 1)$, where we assume that each R_j is a function of $\{O_i : 1 \leq i \leq j\}$. Then,*

$$\frac{1}{n} \sum_{j=0}^{n-1} R_j = o_{P_0}(n^{-\beta}).$$

Conditions C2) through C4) are now easily handled. Condition C2) is a consequence of the fact that

$$\frac{1}{n} \sum_{j=0}^{n-1} \left| \frac{\sigma_j^2}{\hat{\sigma}_j^2} - 1 \right| \leq \gamma^{-1} \frac{1}{n} \sum_{j=0}^{n-1} |\hat{\sigma}_j^2 - \sigma_j^2| \rightarrow 0 \text{ in probability as } n \rightarrow \infty,$$

where the inequality holds by C7) and the convergence holds by C9) Lemma 3. Condition C9) is easily shown to hold under Glivenko-Cantelli conditions on the estimators \hat{P}_j and d_j [see, e.g., Theorem 7 in Luedtke and van der Laan, 2016]. Conditions C3) and C4) are an immediate consequence of C10) and C11) combined with Lemma 3.

While sufficient conditions for C11) should be developed in each individual example, we can give intuition as to why this condition should be reasonable. For any $P \in \mathcal{M}$, let $d(P)$ return a maximizer of (1). We are interested in ensuring that $\Psi^{d_n}(P_0) - \Psi^{d(P_0)}(P_0)$ is small, where d_n is our estimate of a maximizer of (1). This can be expected to hold when the parameter $P \mapsto \Psi^{d(P)}(P_0)$ has pathwise derivative zero at $P = P_0$, where the P_0 in the Ψ argument is fixed. When well-defined, the pathwise derivative will be zero because $d(P)$ is chosen to maximize $\Psi^d(P_0)$ in d .

2.6 Computationally efficient implementation

There are several computationally efficient ways to compute our estimator. In Section 6.1 of Luedtke and van der Laan [2016], we show that the runtime of our estimator can be dramatically improved by running the algorithm used to compute each \hat{P}_j a limited number of times, say ten times. We do not detail this approach here, though we note that the theorems we have presented are general enough to apply to this case.

An alternative approach to improve runtime is to use the estimator’s online nature to compute it efficiently both in time and storage. Suppose that we have an algorithm to update the estimate \hat{P}_{nj} of P_0 to the estimate $\hat{P}_{n(j+1)}$ based on the first j observations by looking at O_{j+1} only. This will often be feasible if the parameter of interest and the bias correction step only require estimates of certain components of P_0 , e.g. of a set of regression and classification functions. In these cases we can apply modern regression and classification approaches to estimate these quantities [see, e.g., Xu, 2011, Luts et al., 2014]. Often d_{nj} can also be obtained using online methods, and thus $\frac{1}{n-\ell_n} \sum_{j=\ell_n}^{n-1} \left[\Psi^{d_{nj}}(\hat{P}_{nj}) + \hat{D}_{nj}(O_{j+1}) \right]$ can be estimated online by keeping a running sum. This quantity is not equal to ψ_n because it does not yet include the weights.

It will not in general be possible to compute the weights online, though their computation does not require storing $O(n)$ observations in memory. We can estimate $\text{Var}_{P_0}(\hat{D}_{nj}(O))$ consistently using the r_j observations, where $r_j \rightarrow \infty$ but can grow very slowly (even $\log j$ suffices asymptotically, though such a slow growth is not recommended for finite samples). Given online estimates of these variances, it is then straightforward to compute both $\bar{\sigma}_n$ and the weights and incorporate these into our estimator. In some cases, we can compute the weights, and thus the estimator, in a truly online fashion. Describing general sufficient conditions for this appears to be difficult, but we conjecture that often this will not typically hold if \mathcal{D}_n is not of finite cardinality. The weights can be computed online in the maximal correlation example that we describe in the next section.

3 Maximal correlation example

3.1 Problem formulation

We now present the running example of this work, namely the maximal correlation estimation problem considered by McKeague and Qian [2015]. The observed data structure is $O = (X, Y)$, where $X = (X_k : k = 1, \dots)$ is a $[-1, 1]^\infty$ vector of predictors and Y is an outcome in $[-1, 1]$. For each n , we let \mathcal{K}_n represent a subset of these predictors of size p , where throughout we assume that

$$\beta_n^2 \equiv \frac{\log p}{\sqrt{n}} \rightarrow 0 \text{ as } n \rightarrow \infty. \quad (5)$$

For readability, we omit the dependence of p on n in the notation. Under a distribution P , the maximal absolute correlation of a predictor with Y is given by

$$\Psi_n(P) \equiv \max_{k \in \mathcal{K}_n} |\text{Corr}_P(X_k, Y)|, \quad (6)$$

where $\text{Corr}_P(X_k, Y)$ is the correlation of X_k and Y under P . We wish to develop confidence intervals for $\Psi_n(P_0)$. When a test of $H_0 : \Psi_n(P_0) = 0$ against the complementary alternative, we also wish to establish the behavior of our test against local alternatives as was done in McKeague and Qian [2015].

In contrast to McKeague and Qian [2015], the procedure that we present in this work:

- 1) is proven to work when p grows with sample size at any rate satisfying (5);
- 2) yields confidence intervals for the maximal correlation rather than just a test of the null hypothesis that it is equal to zero,
- 3) allows a non-null the maximizer in (6) to be non-unique;
- 4) is proven to work in a nonparametric model that neither assumes linearity nor homoscedasticity.

While McKeague and Qian argued that 3) is unlikely in practice, having two non-null maximizers be approximately equal may still have finite sample implications for their test in some settings.

We now show that this problem fits in our framework. To satisfy the pathwise differentiability condition, we let $\mathcal{D}_n \equiv \mathcal{K}_n \times \{-1, 1\}$ and, for each $d = (k, m) \in \mathcal{D}_n$,

$$\Psi^d(P) \equiv m \text{Corr}_P(X_k, Y).$$

Note that $\Psi_n(P)$ now takes the form in (1), where we note that the use of m in the definition of Ψ^d serves to ensure that $\Psi_n(P_0)$ represents the correlation with the maximal *absolute value*.

3.2 Differentiability condition

Canonical gradients

For each k , let $s_P^2(X_k) \equiv \text{Var}_P(X_k)$, and likewise for $s_P^2(Y)$. For ease of notation we let $s_0^2(X_k) \equiv s_{P_0}^2(X_k)$, and likewise for $s_0^2(Y)$ and $\text{Corr}_0(X_k, Y)$. An application of the delta

method shows that Ψ^d has canonical gradient $D^d(P)(o)$ given by

$$m \times \left(\frac{(x_k - E_P[X_k])(y - E_P[Y])}{s_P(X_k)s_P(Y)} - \frac{1}{2} \text{Corr}_P(X_k, Y) \left[\frac{(x_k - E_P[X_k])^2}{s_P^2(X_k)} + \frac{(y - E_P[Y])^2}{s_P^2(Y)} \right] \right).$$

In order to ensure that $D^d(P_0)$ is uniformly bounded for all d , we assume throughout that, for some $\delta \in (0, 1]$,

$$\min\{s_0(Y), s_0(X_1), s_0(X_2), \dots\} > \delta.$$

Second-order remainder

Fix $d = (k, m) \in \mathcal{D}_n$ and $P \in \mathcal{M}$. Let $\tilde{\delta} \in (0, 1]$ be some constant such that both $s_P(X_k)$ and $s_P(Y)$ are larger than $\tilde{\delta}$. Lemma A.1 in the appendix proves that

$$\begin{aligned} |\text{Rem}^d(P)| &\leq \tilde{\delta}^{-1} \left(|s_P(X_k)s_P(Y) - s_0(X_k)s_0(Y)| |\text{Corr}_P(X_k, Y) - \text{Corr}_0(X_k, Y)| \right. \\ &\quad + (E_P[X_k] - E_{P_0}[X_k])^2 + (E_P[Y] - E_{P_0}[Y])^2 \\ &\quad \left. + \frac{s_0^2(Y)}{s_P^2(Y)} [s_P(X_k) - s_0(X_k)]^2 + \frac{s_0^2(X_k)}{s_P^2(X_k)} [s_P(Y) - s_0(Y)]^2 \right). \end{aligned} \quad (7)$$

The first term above is small if $s_P(X_k)$, $s_P(Y)$, and $\text{Corr}_P(X_k, Y)$ are close to $s_0(X_k)$, $s_0(Y)$, and $\text{Corr}_0(X_k, Y)$. The middle terms are small if $E_P[X_k]$ and $E_P[Y]$ are close to $E_{P_0}[X_k]$ and $E_{P_0}[Y]$. The final terms are small if $s_P(X_k)$ and $s_P(Y)$ are close to $s_0(X_k)$ and $s_0(Y)$.

Variance of canonical gradients

There is no elegant (and informative) expression for the variance σ_{nj}^2 of $D^{d_{nj}}(P_j)(O)$. Nonetheless, we show in Lemma A.4 of the appendix that our estimates $\hat{\sigma}_{nj}^2$, taken as the sample variance of $D^{d_{nj}}(P_j)(O)$, concentrate tightly about σ_{nj}^2 with high probability when the sample size is large enough. Thus, in practice, one can actually check if σ_{nj}^2 is small by looking at $\hat{\sigma}_{nj}^2$. If P_0 is normal, then this variance is equal to $[1 - \text{Corr}_{P_0}(X_k, Y)^2]^2$, and so is only zero if $\text{Corr}_{P_0}(X_k, Y) = 1$. Though such an elegant expression does not exist for the variance of $D^d(P_0)(O)$ for general distributions, one can still show in general that the variance of $D^d(P_0)$ is equal to zero only if $\text{Corr}_{P_0}(X_k, Y) = 1$. Here we make the slightly stronger assumption that

$$\inf_{n \geq 2} \min_{(k, m) \in \mathcal{K}_n \times \{-1, 1\}} \text{Var}_{P_0}(D^d(P_0)(O)) \geq \gamma > 0. \quad (8)$$

3.3 Our estimator

We will use the estimator presented in Section 2 to estimate $\Psi_n(P_0)$. At each index $j \geq \ell_n$ we use the empirical distribution P_j of the observations O_1, \dots, O_j to estimate P_0 . We estimate σ_{nj}^2 with the variance of $\hat{D}_{nj}(O)$ under P_j .

In the appendix, we detail conditions on ℓ_n which ensure that ℓ_n does not grow too slowly or quickly. For any $\epsilon \in (0, 2)$, one possible choice of ℓ_n that satisfies these conditions is

$$\ell_n = \max \{ (\log \max\{n, p\})^{1+\epsilon}, n \exp(-\beta_n^{-2+\epsilon}) \}. \quad (9)$$

We show that this choice of ℓ_n ensures C1), C2), and C3) in the appendix. By Theorem 1 this establishes the validity of the lower bound of our confidence interval. We can also show that this lower bound is tight up to a term of the order $n^{-1/4}\beta_n$.

Theorem 4 (Tightness of the lower bound). *For any sequence $t_n \rightarrow \infty$, $\Psi_n(P_0) < \text{LB}_n + t_n n^{-1/4}\beta_n$ with probability approaching 1.*

We now consider the validity of the upper bound of our confidence interval, which holds under C4). This condition is trivially valid if $\Psi_n(P_0) = 0$ for all n . Condition C4) is also valid under the following margin condition:

MC) For some sequence $t_n \rightarrow \infty$, there exists a sequence of non-empty subsets $\mathcal{K}_n^* \subseteq \mathcal{K}_n$ such that, for all n ,

$$\begin{aligned} \sup_{k \in \mathcal{K}_n^*} |\text{Corr}_{P_0}(X_k, Y)| - \inf_{k \in \mathcal{K}_n^*} |\text{Corr}_{P_0}(X_k, Y)| &= o(n^{-1/2}), \\ \inf_{k \in \mathcal{K}_n^*} |\text{Corr}_{P_0}(X_k, Y)| &\geq \sup_{k \in \mathcal{K}_n \setminus \mathcal{K}_n^*} |\text{Corr}_{P_0}(X_k, Y)| + t_n n^{-1/4}\beta_n. \end{aligned}$$

If $\mathcal{K}_n^* = \mathcal{K}_n$, then the supremum over $\mathcal{K}_n \setminus \mathcal{K}_n^*$ is taken to be zero.

Theorem 5 (Validity of the upper bound). *If MC) or $\Psi_n(P_0) = 0$ for all n , then C4) holds so that $\text{LB}_n \leq \Psi_n(P_0) \leq \text{UB}_n$ with probability approaching $1 - \alpha$.*

We outline the techniques used to prove these two results at the end of this subsection. Complete proofs are given in the appendix.

Suppose we wish to test $H_0 : \Psi_n(P_0) = 0$ against $H_1 : \Psi_n(P_0) > 0$. Consider the test that rejects H_0 if $\text{LB}_n > 0$. We wish to explore the behavior of this test under local alternatives

where $\Psi_n(P_0)$ converges to zero slower than $n^{-1/4}\beta_n$. Theorem 4 shows that this test has power converging to one under such local alternatives. Furthermore, as the lower bound is valid in general, this test has type I error of at most $\alpha/2$ under the null. This is indeed an exciting result as it enables the study of local alternatives even when dimension grows quickly with sample size. If dimension does not grow with sample size, this shows that we can detect against any alternatives converging to zero slower than $n^{-1/2}\sqrt{\log n}$. We would not be surprised if the $\sqrt{\log n}$ is unnecessary, but rather that it is simply a result of our proof techniques which give high probability bounds on the concentration of our correlation estimates at each sample size. McKeague and Qian [2015] showed that their method is consistent against a class of alternatives converging to zero slower than $n^{-1/2}$ provided the optimal index is unique. Our result does not rely on this uniqueness condition.

Theorem 5 shows that the upper bound of our confidence interval is also valid under a reasonable margin condition. The margin condition states that there may be many non-null approximate maximizers provided their absolute correlations are well-separated from the absolute correlations of the other predictors with Y . By “approximate” we mean that their absolute correlations all fall within $o(n^{-1/2})$ of one another. If \mathcal{K}_n does not depend on sample size, then this theorem shows that our two-sided confidence interval is always valid.

Sketch of proofs of Theorems 4 and 5. Our proofs of both of these theorems rely on high-probability bounds of the absolute differences between our estimates of $s_{P_j}^2(X_k)$, $s_{P_j}^2(Y)$, $\text{Corr}_{P_j}(X_k, Y)$, and $\hat{\sigma}_{nj}$ and their population counterparts, uniformly over $k \in \mathcal{K}_n$ and j . We show that, with probability at most $1 - 1/n$, all of these absolute differences are upper bounded by constants (with explicit dependence on γ and δ) times $j^{-1/2} \log \max\{n, p\}$.

Condition C1) follows once we show that, with high probability, $s_{P_j}^2(X_k)$ and $s_{P_j}^2(Y)$ are bounded below by $\delta/2$ and $\hat{\sigma}_{nj}^2$ is bounded below by $\gamma/2$ uniformly over $j \geq \ell_n$ for n large enough. Condition C2) and C3) are easy consequences of our concentration results. The concentration results also yield that

$$\frac{1}{n - \ell_n} \sum_{j=\ell_n}^{n-1} \hat{\sigma}_{nj}^{-1} [\Psi^{d_{nj}}(P_0) - \Psi_n(P_0)] = O_{P_0}(n^{-1/4}\beta_n),$$

which then quickly yields Theorem 4 thanks to the expression in (3).

Now suppose MC) holds. By our concentration inequalities, we select a $k_{nj} \in \mathcal{K}_n^*$ for each $j \geq Ct_n^{-1}n$ with high probability, where C is a constant. We also correctly specify m_{nj}

to be the sign of $\text{Corr}_0(X_{k_{nj}}, Y)$. Because all of the absolute correlations in \mathcal{K}_n^* are small, the difference between $\Psi^{d_{nj}}(P_0)$ for $d_{nj} = (k_{nj}, m_{nj})$ and $\Psi_n(P_0)$ is very small. If $\ell_n < Ct_n^{-1}$, then we can apply our concentration inequalities to establish that these first few values of j for which $j < Ct_n^{-1}$ are small enough so that C4) still holds, yielding Theorem 5. \square

In Appendix B, we show that our estimator runs in $O(np)$ time. We show that the estimator can be computed using $O(p)$ storage when the observations O_1, \dots, O_n arrive in a data stream. This result is closely related to the fact that, for a \mathbb{R}^p -valued sequence $\{t_i\}$, the sum $S_j \equiv \sum_{i=1}^j t_i$ at $j = n$ can be computed in time $O(np)$ using storage $O(p)$. In particular, one can use the recursion relation $S_j = t_j + S_{j-1}$, thereby only storing t_j and S_{j-1} when computing S_j . Our estimator can also be computed in $O(np)$ time and $O(n)$ storage when the vectors $(X_{jr} : j = 1, \dots, n) \in \mathbb{R}^n$ arrive in a stream for $r = 1, 2, \dots, p$, where X_{jr} is the observation of X_r for individual j . We do not prove the $O(n)$ storage result in the appendix due to space constraints, though the algorithm is closely related to that given in Appendix B.

4 Simulation study

We now consider the power and scalability of our method using the simulations similar to those described in McKeague and Qian [2015]. Let $X \sim \text{MVN}(0, \Sigma)$ for Σ a $p \times p$ covariance matrix to be given shortly, and τ_1, \dots, τ_p be a sequence of i.i.d. normal random variables independent of all other quantities under consideration. We will use two types of errors: the homoscedastic error τ_1 and the heteroscedastic error $\eta(X) \equiv \sum_{k=1}^p X_k \tau_k / \sqrt{p}$. For $(n, p) = (200, 200), (500, 2000)$, we generate data using the following distributions: (N.IE) $Y = \tau_1$, (A1.IE) $Y = X_1/5 + \tau_1$, (A2.IE) $Y = 0.15 \sum_{k=1}^5 X_k - 0.1 \sum_{k=6}^{10} X_k + \tau_1$, (N.DE) $Y = \eta(X)$, (A1.DE) $Y = X_1/5 + \eta(X)$, and (A2.DE) $Y = 0.15 \sum_{k=1}^5 X_k - 0.1 \sum_{k=6}^{10} X_k + \eta(X)$. For $(n, p) = (2000, 30000)$, we generate data using the following distributions: (N.IE) $Y = \tau_1$, (A3.IE) $Y = X_1/15 + \tau_1$, and (A4.IE) $Y = 0.03 \sum_{k=1}^5 X_k - 0.015 \sum_{k=6}^{10} X_k + \tau_1$. We set all of the diagonal elements in the covariance matrix Σ equal 1, and the off-diagonal elements equal ρ , where for each simulation setting we let $\rho = 0, 0.25, 0.5, 0.75$.

We conduct a 5% test of $\Psi(P_0) > 0$ by checking if the lower bound of a 90% confidence interval for this quantity is greater than zero. We use models N.IE and N.DE to evaluate

type I error and all other models evaluate power. We run our method with ℓ_n as in (9), where we let $\epsilon = 0.5$. For ease of implementation, we compute our method on chunks of data of size $(n - \ell_n)/10$ (see Section 6.1 of Luedtke and van der Laan, 2016). We compare our method to the parametric bootstrap analogue of ART described in Section 2 of Zhang and Laber [2015] for all n . The parametric bootstrap analogue of ART assumes a locally linear model with homoscedastic errors. We use 500 bootstrap draws for each run of the parametric bootstrap procedure. Zhang and Laber show that their method, which does not involve running a computationally burdensome double bootstrap procedure, has comparable performance to ART across sample sizes and predictor dimension, while being more computationally efficient. For this reason, we do not directly compare against the ART results in McKeague and Qian [2015] due to its heavy computational requirements. The parametric bootstrap analogue to ART is less computationally intensive than the ART, but still requires estimating the $p \times p$ covariance matrix Σ and simulating from a $N(0, \hat{\Sigma})$ distribution. Due to computational constraints, we only run this parametric bootstrap analogue for $p \leq 2000$ and not for $p = 30\,000$. We also compare our method to a Bonferroni-corrected t -test.

All simulations are run using 1 000 Monte Carlo simulations in R [R Core Team, 2014].

Figures 1 displays the power of the three testing procedures for (n, p) equal to $(200, 200)$ and $(500, 2000)$ for the homoscedastic data generating distributions N.IE, A1.IE, and A2.IE. The parametric bootstrap analogue of ART performs best in both of these settings. We can show (details omitted) that our method underperforms in this setting due to the second-order term representing the cost for estimating d_0 on subsets of the data of size $j \ll n$ early on in the procedure. While Theorem A.9 ensures that the estimate of d_0 will be asymptotically valid, there appears to be a noticeable price to pay at small sample sizes.

Figures 2 displays the power of the three testing procedures for (n, p) equal to $(200, 200)$ and $(500, 2000)$ for the heteroscedastic data generating distributions. The parametric bootstrap analogue of ART fails to control the type I error in this setting. This is unsurprising given that this test was developed under a local linear model with independent errors. Both our method and Bonferroni adequately control type I error in this setting, especially at the larger sample size $n = 500$, while we see that the Bonferroni procedure achieves slightly better power than our method for these data generating distributions.

Figure 3 displays the power of our method and the Bonferroni procedure for (n, p) equal




Figure 1: Power of the various testing procedures for (n, p) equal to $(200, 200)$ and $(500, 2000)$ under homoscedastic errors. The parametric bootstrap analogue of ART performs the best in this setting.

to $(2000, 30000)$. While (unsurprisingly) Bonferroni performs well when the correlation between the predictors in X is low, our method outperforms the Bonferroni procedure when the correlation increases. We expect that, were we able to run the parametric bootstrap analogue of the ART at this sample size, it would outperform all other methods under consideration as it did at the smaller sample sizes. Nonetheless, this method quickly becomes computationally impractical when p gets large, whereas our procedure and the Bonferroni procedure can still be implemented at these sample sizes. Furthermore, our method is robust to heteroscedastic errors and non-linear data generating distributions.

5 Discussion

We have presented a general method for estimating the (possibly non-unique) maximum of a family of parameter values indexed by $d \in \mathcal{D}_n$. Such an estimation problem is generally non-regular because minor fluctuations of the data generating distribution can change the subset of \mathcal{D}_n for which the corresponding parameter is maximized. Our estimator takes the




Figure 2: Power of the various testing procedures for (n, p) equal to $(200, 200)$ and $(500, 2000)$ under heteroscedastic errors. The parametric bootstrap analogue of ART fails to control the type I error in this setting.

form of a sum of the terms of a martingale difference sequence, which quickly allows us to apply the relevant central limit theorem to study its asymptotics and develop Wald-type confidence intervals. The estimator adapts to the non-regularity of the problem, in the sense that we can give reasonable conditions under which it is regular and asymptotically linear when the maximizer is unique so that regularity is possible.

We have applied our approach to the example of McKeague and Qian [2015] in which one wishes to learn about the maximal absolute correlation between a prespecified outcome and a predictor belonging to some set. The sample splitting that is built into our estimator has enabled us to analyze the estimator when the dimension p of the predictor grows with sample size slowly enough so that $n^{-1/2} \log p \rightarrow 0$ as n goes to infinity. While McKeague and Qian focus on testing the null hypothesis that this maximal absolute correlation is zero, we have established valid confidence intervals for this quantity. The lower bound of our confidence interval is particularly interesting because it is valid under minimal conditions. When p is very large, one might expect that the null of no correlation between the outcome and any of the predictors is unlikely to be true. In these problems, having an estimate of the maximal




Figure 3: Power of the test from the stabilized one-step and from the Bonferroni-adjusted t -test for $(n, p) = (2000, 30\,000)$ under homoscedastic errors. Unsurprisingly Bonferroni performs well when $\rho = 0$. Our method outperforms the Bonferroni procedure when ρ increases.

absolute correlation, or at least a lower bound for this quantity, will likely still be interesting as a measure of the overall relationship between X and Y .

We have also studied the behavior of this null hypothesis test under local alternatives, showing that our test is consistent when the maximal absolute correlation shrinks to zero slower than $n^{-1/2}(\log \max\{n, p\})^{1/2}$. When the dimension of the predictor is fixed, the test of McKeague and Qian is consistent against alternatives shrinking to zero more slowly than $n^{-1/2}$ rather than $(\log n)^{1/2}n^{-1/2}$. We would not be surprised to find that this $(\log n)^{1/2}$ is unnecessary for p fixed and can be removed using more refined proof techniques.

McKeague and Qian do not require that Y and the coordinates of X have range in $[-1, 1]$. We have made this boundedness assumption out of convenience for our proofs and expect that we can replace the boundedness assumptions with appropriate moment assumptions without significantly changing the results. Our simulation results support this claim. The boundedness condition is not as restrictive as it may first seem, as unbounded X and Y can be rescaled to be bounded. Since the sharp null $H_0 : \Psi_n(P_0) = 0$ is invariant to strictly

monotonic transformations of X and Y , our theoretical results yield a valid of H_0 test after applying, e.g., the sigmoid transformation to X and Y .

We note that, in our simulations, the parametric bootstrap analogue of the ART achieves the highest power among competing methods in settings where we are able to run this procedure and the data generating distribution has homoscedastic errors. Nonetheless, this method is invalid under heteroscedastic errors, as we showed in our simulation. The theory for this method is also developed under a linear model, which will not exactly hold in practice. Furthermore, this procedure as currently described is computationally expensive and does not scale well to large data sets, especially when the dimension of the predictor p is large. This difficulty occurs because the procedure requires the computation of a $p \times p$ covariance matrix. The earlier ART method presented in McKeague and Qian [2015], which achieves similar power to its parametric bootstrap analogue, has been shown to be even more computationally burdensome due to its use of a double bootstrap. Thus we believe our method represents an important contribution to the variable screening literature: it is computationally efficient, and has asymptotic theory supporting its power against local alternatives and increasing covariate dimension. None of the earlier works have given rigorous asymptotic theory when the dimension increases with sample size. Given our simulations, we also believe that developing rigorous asymptotic theory under increasing dimension, heteroscedastic errors, and nonlinear data generating distributions for the ART methods is an important area of future work.

The stabilized one-step estimator presented in this paper applies to many other situations not considered in this paper. In an earlier work, we showed that this estimator is useful for estimating the mean outcome under an optimal individualized treatment strategy Luedtke and van der Laan [2016], where the class \mathcal{D}_n now indexes functions mapping from the covariate space to the set of possible treatment decisions. Thanks to the martingale structure of our estimator, the stabilized one-step estimator can be used to construct confidence intervals when the data is drawn sequentially so that the data generating distribution for observation j can depend on that of the first $j - 1$ observations. One interesting example along these lines is to obtain inference for the value of the optimal arm in a multi-armed bandit problem, even in the case where the optimal arm is non-unique and the reward distributions for the optimal arms have different variances. We look forward to seeing further applications of the

general template for a stabilized one-step estimator that we have presented in this paper.

References

- P J Bickel, C A J Klaassen, Y Ritov, and J A Wellner. *Efficient and adaptive estimation for semiparametric models*. Johns Hopkins University Press, Baltimore, 1993.
- B Chakraborty and E E Moodie. *Statistical Methods for Dynamic Treatment Regimes*. Springer, Berlin Heidelberg New York, 2013.
- P Gaenssler, J Strobel, and W Stute. On central limit theorems for martingale triangular arrays. *Acta Math. Hungar.*, 31(3):205–216, 1978.
- K Hirano and J R Porter. Impossibility results for nondifferentiable functionals. *Econometrica*, 80(4):1769–1790, 2012.
- E Laber and S Murphy. Adaptive confidence intervals for the test error in classification. *J. Am. Stat. Assoc.*, 106:904–913, 2011.
- A R Luedtke and M J van der Laan. Statistical inference for the mean outcome under a possibly non-unique optimal treatment strategy. *Ann. Statist.*, 44(2):713–742, 2016.
- J Luts, T Broderick, and M P Wand. Real-time semiparametric regression. *J. Comput. Graph. Statist.*, 23(3):589–615, 2014.
- C McDiarmid. On the method of bounded differences. *Surv. Comb.*, 141(1):148–188, 1989.
- I W McKeague and M Qian. An adaptive resampling test for detecting the presence of significant predictors. *J. Am. Stat. Assoc.*, 110(512), 2015.
- J Pfanzagl. *Estimation in semiparametric models*. Springer, Berlin Heidelberg New York, 1990.
- R Core Team. *R: a language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria, 2014. URL <http://www.r-project.org/>.

M J van der Laan and S D Lendle. Online Targeted Learning. Technical Report 330, available at <http://www.bepress.com/ucbbiostat/>, Division of Biostatistics, University of California, Berkeley, 2014.

M J van der Laan and J M Robins. *Unified methods for censored longitudinal data and causality*. Springer, New York Berlin Heidelberg, 2003.

A W van der Vaart. On differentiable functionals. *Ann. Statist.*, 19:178–204, 1991.

A W van der Vaart and J A Wellner. *Weak convergence and empirical processes*. Springer, Berlin Heidelberg New York, 1996.

B P Welford. Note on a method for calculating corrected sums of squares and products. *Technometrics*, 4(3):419–420, 1962.

W Xu. Towards optimal one pass large scale learning with averaged stochastic gradient descent. *arXiv Prepr. arXiv1107.2490*, 2011.

Y Zhang and E B Laber. Comment. *J Am Stat Assoc*, 110(512):1451–1454, 2015.

Appendix

A Proofs and results for the McKeague and Qian [2015] example

Lemma A.1. Fix $\tilde{\delta} > 0$ and $d \in \mathcal{D}_n$. For any P with $\min_{k \in \mathcal{K}_n} s_P^2(X_k) > \tilde{\delta}$ and $s_P^2(Y) > \tilde{\delta}$, (7) holds.

Proof. Straightforward but tedious calculations show that

$$\begin{aligned} \text{Rem}^d(P) = m \left(\frac{1}{s_P(X_k)s_P(Y)} [s_P(X_k)s_P(Y) - s_0(X_k)s_0(Y)] [\text{Corr}_P(X_k, Y) - \text{Corr}_0(X_k, Y)] \right. \\ \left. + \frac{(E_P[X_k] - E_{P_0}[X_k])(E_P[Y] - E_{P_0}[Y])}{s_P(X_k)s_P(Y)} \right. \\ \left. - \frac{\text{Corr}_P(X_k, Y)}{2} \left[\frac{(E_P[X_k] - E_{P_0}[X_k])^2}{s_P^2(X_k)} + \frac{(E_P[Y] - E_{P_0}[Y])^2}{s_P^2(Y)} \right] \right) \end{aligned}$$

$$- \frac{\text{Corr}_P(X_k, Y)}{2s_P^2(X_k)s_P^2(Y)} [s_P(X_k)s_0(Y) - s_0(X_k)s_P(Y)]^2 \Big). \quad (\text{A.1})$$

The result follows by taking the absolute value of both sides, applying the triangle inequality, using that $ab \leq (a^2 + b^2)/2$ for any real a, b , $\text{Corr}_P(X_k, Y) \leq 1$, and the lower bound $\tilde{\delta}$ on the variances. \square

We now establish high probability bounds on the difference between $s_{P_j}^2(X_k)$, $s_{P_j}^2(Y)$, $\text{Corr}_{P_j}(X_k, Y)$, and $\hat{\sigma}_{nj}$ and their population counterparts, uniformly over $k \in \mathcal{K}_n$ and j . We will use \lesssim to denote “less than or equal to up to a universal multiplicative constant”. Let \mathcal{F}_n denote the following class of functions mapping from $\mathcal{O} \equiv \mathbb{R}^\infty \times \mathbb{R}$ to the real line:

$$\{(x, y) \mapsto x_k^r y^s : 0 \leq r, s, \leq 4; r + s \leq 4 \text{ } k \in \mathcal{K}_n\}. \quad (\text{A.2})$$

Note that $|\mathcal{F}_n| \lesssim p$. We will use this class to develop concentration results about our estimates the needed portions of the likelihood. This class is actually somewhat larger than is needed for most of our results, as in fact

$$\begin{aligned} & \{(x, y) \mapsto x_k y : k \in \mathcal{K}_n\} \cup \{(x, y) \mapsto x_k : k \in \mathcal{K}_n\} \cup \{(x, y) \mapsto y\} \\ & \cup \{(x, y) \mapsto x_k^2 : k \in \mathcal{K}_n\} \cup \{(x, y) \mapsto y^2\} \end{aligned}$$

suffices for concentrating our estimates of $\text{Corr}_0(X_k, Y)$, $s_0(X_k)$, and $s_0(Y)$. Nonetheless, using this larger class \mathcal{F}_n will allow us to prove results about the concentration of $\hat{\sigma}_{nj}^2$ about σ_{nj}^2 , and just stating it as a single class is convenient for brevity.

For $f \in \mathcal{F}_n$ and $j \in \{1, \dots, n\}$, define the empirical process as

$$\mathbb{G}_{nj} \equiv \frac{1}{\sqrt{j}} \sum_{i=1}^j [f(O_i) - P_0 f] = \sqrt{j}(P_j - P_0)f,$$

where we use P_j denote the empirical distribution of O_0, \dots, O_{j-1} and $Pf \equiv E_P[f(O)]$ for any distribution P . Let $\|\mathbb{G}_{nj}\|_{\mathcal{F}_n} \equiv \sup_{f \in \mathcal{F}_n} |\mathbb{G}_{nj}|$. By Theorem 2.14.1 in van der Vaart and Wellner [1996] shows that

$$E \|\mathbb{G}_{nj}\|_{\mathcal{F}_n} \lesssim \sqrt{\log \#\mathcal{F}_n} \lesssim \sqrt{\log p}, \quad (\text{A.3})$$

where the expectation is over the draws O_1, \dots, O_j . We have used that our class is bounded by the constant 1.

Let

$$K_{nj} \equiv j^{-1/2} \sqrt{\log \max\{n, p\}}. \quad (\text{A.4})$$

Define the events

$$\begin{aligned} \mathcal{A}_{nj} &\equiv \left\{ \max_{f \in \mathcal{F}_n} |(P_j - P_0)f| \leq CK_{nj} \right\} \text{ for all } j = 1, \dots, n, \\ \mathcal{A}_n &\equiv \cap_{j=1}^n \mathcal{A}_{nj}, \end{aligned}$$

where C in the definition of \mathcal{A}_{nj} is equal the smallest universal constant satisfying (A.3) plus 1.

Lemma A.2. *For any sample size n , the event \mathcal{A}_n occurs with probability at least $1 - n/\max\{n^2, p\} \geq 1 - 1/n$.*

Proof. We first upper bound the probability of the complement of \mathcal{A}_{nj} for each n, j . Fix n and $j \leq n$. By the bounds on X and Y , changing one O_i in (O_1, \dots, O_j) to some other value in the support of P_0 can change b by at most $1/\sqrt{j}$. Thus $(O_1, \dots, O_j) \mapsto \|\mathbb{G}_{nj}\|_{\mathcal{F}_n}$ satisfies the bounded differences property with bound $1/\sqrt{j}$, and we may apply McDiarmid's inequality [McDiarmid, 1989] to show that, with probability at most $1 - \exp(-2t^2)$, $\|\mathbb{G}_{nj}\|_{\mathcal{F}_n} \leq E \|\mathbb{G}_{nj}\|_{\mathcal{F}_n} + t$. Choosing $t = \sqrt{\frac{\log \max\{n^2, p\}}{2}}$ and using (A.3) yields that, with probability at least $1 - 1/\max\{n^2, p\}$, the following inequality holds for all $j = 1, \dots, n$:

$$\begin{aligned} \|\mathbb{G}_{nj}\|_{\mathcal{F}_n} &\leq E \|\mathbb{G}_{nj}\|_{\mathcal{F}_n} + \sqrt{\frac{\log \max\{n^2, p\}}{2}} \leq C' \sqrt{\log p} + \sqrt{\log \max\{n, p\}} \\ &\leq C \sqrt{\log \max\{n, p\}}, \end{aligned}$$

where C' denotes the universal constant in (A.3).

By DeMorgan's laws and a union bound, it follows that the event $\mathcal{A}_n \equiv \cap_j \mathcal{A}_{nj}$ occurs with probability at least $1 - n/\max\{n^2, p\} \geq 1 - 1/n$. \square

We have shown that \mathcal{A}_n occurs with high probability. Now we show that our estimates of variances, covariances, and correlations perform well when \mathcal{A}_n occurs.

Lemma A.3. *Fix a sample size $n \geq 2$. The occurrence of \mathcal{A}_n implies that, for all $j = 2, \dots, n$:*

- 1) $\max_{k \in \mathcal{K}_n} |s_{P_j}(X_k) - s_0(X_k)| \lesssim \delta^{-1/2} K_{nj};$
- 2) $\max_{k \in \mathcal{K}_n} |s_{P_j}^2(X_k) - s_0^2(X_k)| \lesssim K_{nj};$
- 3) $|s_{P_j}(Y) - s_0(Y)| \lesssim \delta^{-1/2} K_{nj};$
- 4) $|s_{P_j}^2(Y) - s_0^2(Y)| \lesssim K_{nj};$
- 5) $\max_{k \in \mathcal{K}_n} |\text{Corr}_{P_j}(X_k, Y) - \text{Corr}_{P_0}(X_k, Y)| \lesssim \delta^{-1} K_{nj},$

where we define $\text{Corr}_{P_j}(X_k, Y) = 0$ when either $s_{P_j}(X_k)$ or $s_{P_j}(Y)$ is equal to zero.

Proof. Suppose \mathcal{A}_n holds and fix $k \in \mathcal{K}_n$. The triangle inequality and the bounds on X_k yield that

$$\begin{aligned} |s_{P_j}^2(X_k) - s_0^2(X_k)| &= |(E_{P_j}[X_k^2] - E_{P_0}[X_k^2]) - (E_{P_j}[X_k] + E_{P_0}[X_k]) (E_{P_j}[X_k] - E_{P_0}[X_k])| \\ &\leq |E_{P_j}[X_k^2] - E_{P_0}[X_k^2]| + 2 |E_{P_j}[X_k] - E_{P_0}[X_k]| \lesssim K_{nj}. \end{aligned}$$

This gives 2). For 1), note that

$$|s_{P_j}(X_k) - s_0(X_k)| = \left| \frac{s_{P_j}^2(X_k) - s_0^2(X_k)}{s_{P_j}(X_k) + s_0(X_k)} \right| \lesssim \delta^{-1/2} K_{nj}.$$

The same argument yields 3) and 4).

Again fix k . An application of the triangle inequality and the bounds on X_k and Y readily yield that $|\text{Cov}_{P_j}(X_k, Y) - \text{Cov}_{P_0}(X_k, Y)| \lesssim K_{nj}$. Furthermore,

$$\begin{aligned} &\text{Corr}_{P_j}(X_k, Y) - \text{Corr}_{P_0}(X_k, Y) \\ &= \frac{\text{Cov}_{P_j}(X_k, Y) - \text{Cov}_{P_0}(X_k, Y)}{s_0(X_k)s_0(Y)} - \frac{\text{Corr}_{P_j}(X_k, Y)s_{P_j}(Y)}{s_0(X_k)s_0(Y)} [s_{P_j}(X_k) - s_0(X_k)] \\ &\quad - \frac{\text{Corr}_{P_j}(X_k, Y)}{s_0(Y)} [s_{P_j}(Y) - s_0(Y)]. \end{aligned}$$

Taking the absolute value of both sides, applying the triangle inequality, and using the lower bounds on $s_0(X_k)$ and $s_0(Y)$ and the upper bound on $\text{Corr}_{P_j}(X_k, Y)$ yields that $|\text{Corr}_{P_j}(X_k, Y) - \text{Corr}_{P_0}(X_k, Y)| \lesssim \delta^{-1} K_{nj}$. This holds for all k , so 5) holds. \square

Lemma A.4. *Let C be the smallest universal constant in 2) of that Lemma A.3, and let n be any natural number satisfying $n \geq \lceil 4C^2\delta^{-2} \log \max\{n, p\} \rceil \equiv J(n, \delta)$. Under these conditions, the occurrence of \mathcal{A}_n implies that, for all $j = J(n, \delta), \dots, n$,*

$$6) \min_{k \in \mathcal{K}_n} s_{P_j}^2(X_k) \geq \delta/2 \text{ and } \min_{k \in \mathcal{K}_n} s_{P_j}^2(Y) \geq \delta/2;$$

- 7) $\max_{k \in \mathcal{K}_n} \frac{s_{P_0}^2(X_k)}{s_{P_j}^2(X_k)} \leq 2$ and $\frac{s_{P_0}^2(Y)}{s_{P_j}^2(Y)} \leq 2$;
8) $|\hat{\sigma}_{nj}^2 - \sigma_{nj}^2| \lesssim \delta^{-2} K_{nj}$;
9) $|\hat{\sigma}_{nj}^2 - \text{Var}_{P_0}(D_{nj}(P_0)(O))| \lesssim \delta^{-2} K_{nj} + \left(\widehat{\text{Rem}}_{nj}\right)^2$.

Proof. By Lemma A.3, 2) holds, and using that $j \geq J(n, \delta)$, we see that

$$s_{P_j}^2(X_k) = s_{P_0}^2(X_k) + s_{P_j}^2(X_k) - s_{P_0}^2(X_k) \geq s_{P_0}^2(X_k) - \max_{k \in \mathcal{K}_n} |s_{P_j}^2(X_k) - s_{P_0}^2(X_k)| \geq \delta/2.$$

The same argument works for $s_{P_j}^2(Y)$, so 6) holds. Furthermore,

$$\frac{s_{P_0}^2(X_k)}{s_{P_j}^2(X_k)} \leq \frac{s_{P_0}^2(X_k)}{s_{P_0}^2(X_k) - CK_{nj}} = 1 + \frac{CK_{nj}}{s_{P_0}^2(X_k) - CK_{nj}} \leq 1 + 2C\delta^{-1}K_{nj} \leq 2,$$

where the final two inequalities hold by 6). This proves the first part of 7), and the bound on $s_{P_0}^2(Y)/s_{P_j}^2(Y)$ holds by the same argument. For the second result, note that

$$\begin{aligned} |\hat{\sigma}_{nj}^2 - \sigma_{nj}^2| &\leq |(P_j - P_0)\hat{D}_{nj}^2| + |(P_j\hat{D}_{nj})^2 - (P_0\hat{D}_{nj})^2| \\ &\leq |(P_j - P_0)\hat{D}_{nj}^2| + |(P_j + P_0)\hat{D}_{nj}| |(P_j - P_0)\hat{D}_{nj}|. \end{aligned}$$

Using 6), the bounds on X and Y , and the triangle inequality shows that

$$\lesssim |(P_j - P_0)\hat{D}_{nj}^2| + \delta^{-1} |(P_j - P_0)\hat{D}_{nj}| \lesssim \delta^{-2} \|\mathbb{G}_{nj}\|_{\mathcal{F}_n},$$

where we have used that \mathcal{F}_n contains all polynomials of X_k, Y of degree at most 4. By the occurrence of \mathcal{A}_n , the final line is upper bounded by a constant times $\delta^{-2} K_{nj}$. This yields 8).

For 9), we will bound $|\sigma_{nj}^2 - \text{Var}_{P_0}(D_{nj}(P_0)(O))|$ and then combine this with 8) using the triangle inequality. We have that

$$|\sigma_{nj}^2 - \text{Var}_{P_0}(D_{nj}(P_0)(O))| \leq \left| P_0 \left[\hat{D}_{nj}^2 - D_{nj}(P_0)^2 \right] \right| + \left(P_0 \hat{D}_{nj} \right)^2.$$

Now we use that $P_0 \hat{D}_{nj} = -\text{Corr}_{P_j}(X_k, Y) + \text{Corr}_{P_0}(X_k, Y) + \widehat{\text{Rem}}_{nj}$ and $(a+b)^2 \leq 2(a^2 + b^2)$ for any real a, b to see that $\left(P_0 \hat{D}_{nj} \right)^2 \lesssim \max_k (\text{Corr}_{P_j}(X_k, Y) - \text{Corr}_{P_0}(X_k, Y))^2 + \left(\widehat{\text{Rem}}_{nj} \right)^2$. By 5) from Lemma A.3 and the fact that $j \geq J(n, \delta)$, the maximum over $k \in \mathcal{K}_n$ is bounded above by a constant times $\delta^{-2} K_{nj}^2 \lesssim \delta^{-2} K_{nj}$. Continuing with the above,

$$\lesssim \left| P_0 \left(\left[\hat{D}_{nj} + D_{nj}(P_0) \right] \left[\hat{D}_{nj} - D_{nj}(P_0) \right] \right) \right| + \delta^{-2} K_{nj} + \left(\widehat{\text{Rem}}_{nj} \right)^2$$

$$\begin{aligned}
&\lesssim \delta^{-1} P_0 \left| \hat{D}_{nj} - D_{nj}(P_0) \right| + \delta^{-2} K_{nj} + \left(\widehat{\text{Rem}}_{nj} \right)^2 \\
&\lesssim \delta^{-2} \|\mathbb{G}_{nj}\|_{\mathcal{F}_n} + \delta^{-2} K_{nj} + \left(\widehat{\text{Rem}}_{nj} \right)^2 \\
&\lesssim \delta^{-2} K_{nj} + \left(\widehat{\text{Rem}}_{nj} \right)^2,
\end{aligned}$$

where we used 6) for the second to last inequality. \square

Lemma A.5. *Suppose the conditions of Lemma A.4. Under these conditions, the occurrence of \mathcal{A}_n implies that, for all $j = J(n, \delta), \dots, n$,*

$$10) \left| \widehat{\text{Rem}}_{nj} \right| \lesssim \delta^{-5/2} (K_{nj})^2.$$

Proof. By Lemma A.4, $\min_{k \in \mathcal{K}_n} s_{P_j}^2(X_k) \geq \delta/2$ and $\min_{k \in \mathcal{K}_n} s_{P_j}^2(Y) \geq \delta/2$. By Lemma A.1, this yields

$$\begin{aligned}
\left| \widehat{\text{Rem}}_{nj} \right| &\lesssim \delta^{-1} \max_{k \in \mathcal{K}_n} \left(\left| s_{P_j}(X_k) s_{P_j}(Y) - s_0(X_k) s_0(Y) \right| \left| \text{Corr}_{P_j}(X_k, Y) - \text{Corr}_0(X_k, Y) \right| \right. \\
&\quad + \left(E_{P_j}[X_k] - E_{P_0}[X_k] \right)^2 + \left(E_{P_j}[Y] - E_{P_0}[Y] \right)^2 \\
&\quad \left. + \frac{s_0^2(Y)}{s_{P_j}^2(Y)} [s_{P_j}(X_k) - s_0(X_k)]^2 + \frac{s_0^2(X_k)}{s_{P_j}^2(X_k)} [s_{P_j}(Y) - s_0(Y)]^2 \right).
\end{aligned}$$

By the bounds on X and Y and the triangle inequality, $\left| s_{P_j}(X_k) s_{P_j}(Y) - s_0(X_k) s_0(Y) \right| \leq \left| s_{P_j}(X_k) - s_{P_0}(X_k) \right| + \left| s_{P_j}(Y) - s_{P_0}(Y) \right|$. Applying 7) from Lemma A.4 and the results of Lemma A.3 to the above yields the result. \square

Lemma A.6. *Let γ be as defined in (8). For a constant $C(\gamma, \delta) > 0$ relying on γ and δ only, the occurrence of \mathcal{A}_n implies that, for all $j = \lceil C(\gamma, \delta) \log \max\{n, p\} \rceil, \dots, n$,*

$$11) \hat{\sigma}_{nj}^2 \geq \gamma/2.$$

Sketch of proof. Suppose \mathcal{A}_n . By 9) and 10), for all $j \geq J(n, \delta)$

$$\left| \hat{\sigma}_{nj}^2 - \text{Var}_{P_0}(D_{nj}(P_0)(O)) \right| \lesssim \delta^{-2} K_{nj} + \delta^{-10} (K_{nj})^4.$$

It is easy to confirm that, for a universal constant $C > 0$, the above yields that the left-hand side is upper bounded by $\gamma/2$ for all $j \geq C\gamma^{-1/2}\delta^{-2} \max\{\delta^{-3}, \gamma^{-3/2}\} \log \max\{n, p\} \equiv C(\gamma, \delta) \log \max\{n, p\} \geq J(n, \delta)$. An application of the triangle inequality gives the result. \square

The remainder of the results in this section are asymptotic in nature. We omit the dependence on δ and γ in these statements as these quantities are treated as fixed as sample size grows. Throughout we assume that

$$\frac{\log \max\{n, p\}}{\ell_n} \rightarrow 0, \quad (\text{A.5})$$

$$\beta_n^2 \log \frac{n}{\ell_n} \rightarrow 0, \quad (\text{A.6})$$

$$\limsup_{n \rightarrow \infty} \frac{\ell_n}{n} < 1. \quad (\text{A.7})$$

In view of (A.5) and (A.6), we see that, roughly, ℓ_n grows faster than $\log \max\{n, p\}$ if β_n goes to zero faster than $1/\sqrt{\log n}$ and at least as fast as $n \exp(-o(\beta_n^{-2}))$ if β_n goes to zero more slowly than $1/\sqrt{\log n}$. Given an $\epsilon > 0$, one possible choice of ℓ_n that satisfies these properties is

$$\ell_n = \max \left\{ (\log \max\{n, p\})^{1+\epsilon}, n \exp(-\beta_n^{-2+\epsilon}) \right\}$$

We have the following result.

Lemma A.7. *For all n large enough, $\ell_n \geq J(n, \delta)$ and $\ell_n \geq C(\gamma, \delta) \log \max\{n, p\}$ as defined Lemmas A.4 and A.6, respectively.*

Proof. This is an immediate consequence (A.5) of the fact that δ and γ are fixed as sample size grows. \square

Theorem A.8. *C1), C2), and C3) hold.*

Proof. C1): By Lemma A.7, we can apply 6) from Lemma A.4 and Lemma A.6 provided n is large enough. In that case, $\frac{\hat{D}_{nj}}{\hat{\sigma}_{nj}} \lesssim \delta^{-1} \gamma^{-1/2}$ for all $j \geq \ell_n$ provided \mathcal{A}_n holds. By Lemma A.2, this then occurs with probability at least $1 - 1/n$, and thus C1) holds.

C2): If \mathcal{A}_n holds, then Lemmas A.6 and A.7 show that, for all n large enough,

$$\frac{1}{n - \ell_n} \sum_{j=\ell_n}^{n-1} \left| \frac{\sigma_{nj}^2}{\hat{\sigma}_{nj}^2} - 1 \right| \leq \frac{2\gamma^{-1}}{n - \ell_n} \sum_{j=\ell_n}^{n-1} |\hat{\sigma}_{nj}^2 - \sigma_{nj}^2| \lesssim \frac{\gamma^{-1} \delta^{-2}}{n - \ell_n} \log \max\{n, p\} \sum_{j=\ell_n}^{n-1} j^{-1}$$

By 8) in Lemma A.4 and the fact that $\sum_a^b j^{-1} \leq \int_{a-1}^b j^{-1} dj$, the right-hand side is has an upper bound proportional to $\frac{\gamma^{-1} \delta^{-2}}{n - \ell_n} \log \max\{n, p\} \log n$. This bound is $o(1)$ by (A.7) and the

fact that $\beta_n \rightarrow 0$. The fact that \mathcal{A}_n occurs with probability approaching 1 (Lemma A.2) yields C2).

C3): Suppose that n is large enough so that the results of Lemma A.7 apply. Also suppose that \mathcal{A}_n occurs. We have that

$$\begin{aligned}
\frac{1}{n - \ell_n} \sum_{j=\ell_n}^{n-1} \frac{\widehat{\text{Rem}}_{nj}}{\hat{\sigma}_{nj}} &\lesssim \frac{\gamma^{-1/2} \delta^{-5/2}}{n - \ell_n} \sum_{j=\ell_n}^{n-1} (K_{nj})^2 && \text{(Lemmas A.5, A.6, and A.7)} \\
&= \frac{\gamma^{-1/2} \delta^{-5/2}}{n - \ell_n} \log \max\{n, p\} \sum_{j=\ell_n}^{n-1} j^{-1} && \text{(Eq. A.4)} \\
&\lesssim \frac{\gamma^{-1/2} \delta^{-5/2}}{n - \ell_n} \log \max\{n, p\} \log \frac{n}{\ell_n} && (\sum_a^b j^{-1} \leq \int_{a-1}^b j^{-1} dj) \\
&= o([n - \ell_n]^{-1/2}). && \text{(Eqs. A.6 and A.7)}
\end{aligned}$$

The fact that \mathcal{A}_n occurs with probability approaching 1 (Lemma A.2) yields C3). \square

Let k_{n0} be a possibly non-unique k maximizer of $|\text{Corr}_{P_0}(X_k, Y)|$. For each $r > 0$, let $\mathcal{K}_n^r \subseteq \mathcal{K}_n$ denote the set of all $k \in \mathcal{K}_n$ such that $|\text{Corr}_{P_0}(X_{k_{n0}})| - |\text{Corr}_{P_0}(X_k, Y)| \leq r$.

The upcoming theorem uses the following conditions to establish the validity of a hypothesis test of no effect and of the upper bound of our confidence interval, respectively:

M1) For some sequence $\{t_n\}$ with $t_n \rightarrow +\infty$, there exists a sequence of non-empty subsets $\mathcal{K}_n^* \subseteq \mathcal{K}_n$ such that, for all n ,

$$\inf_{k_1 \in \mathcal{K}_n^*} |\text{Corr}_{P_0}(X_{k_1}, Y)| \geq \sup_{k_2 \in \mathcal{K}_n \setminus \mathcal{K}_n^*} |\text{Corr}_{P_0}(X_{k_2}, Y)| + t_n n^{-1/4} \beta_n.$$

If $\mathcal{K}_n^* = \mathcal{K}_n$, then the supremum on the right-hand side is taken to be zero.

M2) The conditions of M1) hold, and also

$$\text{Diam}(\mathcal{K}_n^*) \equiv \sup_{k_1, k_2 \in \mathcal{K}_n^*} (|\text{Corr}_{P_0}(X_{k_1}, Y)| - |\text{Corr}_{P_0}(X_{k_2}, Y)|) = o(n^{-1/2}).$$

The first of these conditions will be used to establish the consistency of a null hypothesis significance test. The second of these conditions is similar to margin conditions used in classification, and will be used to establish the validity of our confidence interval.

Theorem A.9.

$$\frac{1}{n - \ell_n} \sum_{j=\ell_n}^{n-1} \hat{\sigma}_{nj}^{-1} [\Psi^{d_{nj}}(P_0) - \Psi_n(P_0)] = O_{P_0}(n^{-1/4}\beta_n). \quad (\text{A.8})$$

If also M1), then the right-hand side of the above can be tightened to $O_{P_0}(\text{Diam}(\mathcal{K}_n^*) \wedge n^{-1/4}\beta_n) + o_{P_0}(n^{1/2})$. If also M2), then C4) holds.

Proof. Suppose that \mathcal{A}_n holds and n is large enough so that the results of Lemma A.7 apply. For each $j \geq \ell_n$, let k_{nj} represent the $k \in \mathcal{K}_n$ which maximizes $|\text{Corr}_{P_j}(X_k, Y)|$. Let $m_0 = \text{sgn}[\text{Corr}_{P_0}(X_{k_{n0}}, Y)]$ and $m_{nj} = \text{sgn}[\text{Corr}_{P_j}(X_{k_{nj}}, Y)]$. Then, for a universal constant $C > 0$,

$$\begin{aligned} 0 &\geq m_0 \text{Corr}_{P_j}(X_{k_{n0}}, Y) - |\text{Corr}_{P_j}(X_{k_{nj}}, Y)| \\ &= [m_0 \text{Corr}_{P_0}(X_{k_{n0}}, Y) - m_{nj} \text{Corr}_{P_0}(X_{k_{nj}}, Y)] \\ &\quad + m_0 [\text{Corr}_{P_j}(X_{k_{n0}}, Y) - \text{Corr}_{P_0}(X_{k_{n0}}, Y)] - m_{nj} [\text{Corr}_{P_j}(X_{k_{nj}}, Y) - \text{Corr}_{P_0}(X_{k_{nj}}, Y)] \\ &\geq \Psi_n(P_0) - \Psi^{d_{nj}}(P_0) - 2 \max_{k \in \mathcal{K}_n} |\text{Corr}_{P_j}(X_k, Y) - \text{Corr}_{P_0}(X_k, Y)| \\ &\geq \Psi_n(P_0) - \Psi^{d_{nj}}(P_0) - C\delta^{-1}K_{nj}, \end{aligned} \quad (\text{A.9})$$

where the final inequality holds by Lemma 5). Using that $\sum_{j=\ell_n}^{n-1} j^{-1/2} \lesssim \sqrt{n}$ and (A.7),

$$\frac{1}{n - \ell_n} \sum_{j=\ell_n}^{n-1} K_{nj} \lesssim \log \max\{n, p\} \frac{\sqrt{n}}{n - \ell_n} \lesssim n^{-1/4}\beta_n.$$

By Lemma A.6, this then implies that the left-hand side of (A.8) is upper bounded by an $O(\gamma^{-1/2}n^{-1/4}\beta_n)$ term under \mathcal{A}_n , and so Lemma A.2 yields (A.8).

For the second result, suppose that M1) holds. Observe that, for all $j > Cnt_n^{-1}$ for C as defined in (A.9), $\Psi_n(P_0) - \Psi^{d_{nj}}(P_0) < t_n n^{-1/4}\beta_n$. Furthermore, $\Psi_n(P_0) - |\text{Corr}_{P_0}(X_{k_{nj}})| \leq \Psi_n(P_0) - \Psi^{d_{nj}}(P_0)$. Thus $k_{nj} \in \mathcal{K}_n^*$ as defined in M1). Furthermore, m_{nj} must equal $\text{sgn}[\text{Corr}_{P_0}(X_{k_{nj}}, Y)]$, since otherwise

$$\Psi_n(P_0) - \Psi^{d_{nj}}(P_0) \geq |\text{Corr}_{P_0}(X_{k_{nj}}, Y)| - \Psi^{d_{nj}}(P_0) \geq 2 \inf_{k \in \mathcal{K}_n^*} |\text{Corr}_{P_0}(X_k, Y)| \geq 2t_n n^{-1/4}\beta_n,$$

contradicting the fact that $\Psi_n(P_0) - \Psi^{d_{nj}}(P_0) \leq t_n n^{-1/4}\beta_n$ per (A.9). Because $k \in \mathcal{K}_n^*$, we see that $\Psi^{d_{nj}}(P_0) \geq \inf_{k \in \mathcal{K}_n^*} |\text{Corr}_{P_0}(X_k, Y)|$. Hence,

$$\frac{1}{n - \ell_n} \sum_{j=\max\{\ell_n, \lceil Cnt_n^{-1} \rceil\}}^{n-1} [\Psi^{d_{nj}}(P_0) - \Psi_n(P_0)] \geq -\text{Diam}(\mathcal{K}_n^*). \quad (\text{A.10})$$

Further, if $\lceil Cnt_n^{-1} \rceil \geq \ell_n$, (A.9) yields

$$\sum_{j=\ell_n}^{\lceil Cnt_n^{-1} \rceil} [\Psi^{d_{nj}}(P_0) - \Psi_n(P_0)] \geq -C \sum_{j=\ell_n}^{\lceil Cnt_n^{-1} \rceil} K_{nj} \geq -C \log \max\{n, p\} \int_{\ell_n-1}^{\lceil Cnt_n^{-1} \rceil} j^{-1/2} dj$$

It follows that the left-hand side above is greater than or equal to a positive universal constant times $n^{1/2}t_n^{-1/2}$. Dividing the left by $n - \ell_n$ and applying (A.7) yields that this same result holds with an upper bound on the order of $n^{-1/2}t_n^{-1/2}$. Combining this with (A.10) shows that

$$\frac{1}{n - \ell_n} \sum_{j=\ell_n}^{n-1} [\Psi^{d_{nj}}(P_0) - \Psi_n(P_0)] \geq -\text{Diam}(\mathcal{K}_n^*) + O(n^{-1/2}t_n^{-1/2}).$$

Using that $t_n^{-1/2} \rightarrow 0$, $n^{-1/2}t_n^{-1/2} = o(n^{-1/2})$. When proving the first result (A.8) we also showed that the left-hand side is upper-bounded by a positive constant times $-\delta^{-1}n^{-1/4}\beta_n$. Combining with Lemma A.6 and using that \mathcal{A}_n holds with probability approaching 1 (Lemma A.2) shows that the left-hand side of (A.8) is $O_{P_0}(\text{Diam}(\mathcal{K}_n^*) \wedge n^{-1/4}\beta_n) + o_{P_0}(n^{-1/2})$. If M2) holds, then this expression is $o_{P_0}(n^{-1/2})$, and so C4) holds. \square

B Computationally efficient implementation of our estimator for the McKeague and Qian [2015] example

In this section, we describe how to implement the estimator in $O(np)$ time. We show that this can be accomplished using $O(p)$ storage when the observations O_1, \dots, O_n arrive in a stream.

Fix n so that the set \mathcal{K}_n of predictor indices is also fixed. For each j , let P_j denote the empirical distribution of the first j observations. Recall the definition of the class \mathcal{F}_n from (A.2), and note that \mathcal{F}_n contains $O(p)$ functions. It is easy to see that, at $j = 2$, we can compute $P_j f \equiv E_{P_j}[f(O)]$ for each $f \in \mathcal{F}_n$ using $O(p)$ time and storage. Furthermore, for $j > 3$ the fact that $P_j f = f(O_j) + \frac{j-1}{j}P_{j-1}f$ shows that we can compute and save $P_j f$ in $O(p)$ time and storage if we know O_j and $P_{j-1}f$. To attain this storage complexity, we remove $P_{j-2}f$, $f \in \mathcal{F}_n$, from memory for each $j \geq 4$ so that $P_2f, \dots, P_{j-2}f$ are not stored in memory.

We now have an algorithm that, at observation j , starts with O_j and $P_{j-1}f$, $f \in \mathcal{F}_n$, stored in memory and, after running the steps described in the preceding paragraph, also has P_jf , $f \in \mathcal{F}_n$ stored in memory. Given P_jf , $f \in \mathcal{F}_n$, one can compute and save $\text{Cov}_{P_j}(X_k, Y) = E_{P_j}[X_k Y] - E_{P_j}[X_k]E_{P_j}[Y]$, $k \in \mathcal{K}_n$, and $s_{P_j}^2(Z) = E_{P_j}[Z^2] - E_{P_j}[Z]^2$, Z equal to Y or X_k , $k \in \mathcal{K}_n$, in $O(p)$ time and storage. We can now compute and save $\text{Corr}_{P_j}(X_k, Y) = \frac{\text{Cov}_{P_j}(X_k, Y)}{s_{P_j}(X_k)s_{P_j}(Y)}$, $k \in \mathcal{K}_n$, in $O(p)$ time and storage. If the predictors or outcome are large and their variance small, the described online computation of the sample variance may lead to numerical difficulties. See Welford [1962] for a better estimate of the variance in this setting.

Let H_j denote the collection of (i) the integer j , (ii) P_jf , $f \in \mathcal{F}_n$, (iii) $s_{P_j}^2(Y)$, and (iv) $\text{Cov}(X_k, Y)$, $s_{P_j}^2(X_k)$ and $\text{Corr}_{P_j}(X_k, Y)$, $k \in \mathcal{K}_n$. For $j \geq 2$, let `UPDATEH` be a function which takes as input (O_{j+1}, H_j) and outputs H_{j+1} . We have shown that `UPDATEH` (O_{j+1}, H_j) can run in $O(p)$ time for any $j \geq 2$. We call a separate function `INITIALIZEH` on (O_1, O_2) to obtain the initial value H_2 . This function runs in $O(p)$ time and storage.

Let the function `MAXIMIZER` be a function that takes as input H_j and returns the $d_j = (k_j, m_j)$ which maximizes $m \text{Corr}_{P_j}(X_k, Y)$ in $d \equiv (k, m) \in \mathcal{K}_n$, thereby allowing us to compute $\hat{\sigma}_{nj} = P_j D^{d_j}(P_j)^2$. Finding d_j involves finding the maximum of $|\mathcal{D}_n| = 2p$ numbers, and therefore can be accomplished in $O(p)$ time.

The function `CALCD` takes as input H_j , O_{j+1} , and d_j and calculates $D^{d_j}(P_j)(O_{j+1})$. It is easy to see that this can be accomplished in $O(1)$ time and $O(p)$ storage.

For ease of notation in the proceeding paragraph and equation we omit the dependence of $d_j = (k_j, m_j)$ on j in the notation. Since $D^d(P_j)$ is a gradient for Ψ^d at P_j and gradients are mean zero, $P_j D^d(P_j) = 0$. For any $d \in \mathcal{D}_n$, tedious but trivial calculations show that

$$\begin{aligned} P_j D^d(P_j)^2 &= \left[\frac{2 + \text{Corr}_{P_j}(X_k, Y)^2}{2s_{P_j}^2(X_k)s_{P_j}^2(Y)} \right] \sum_{r=0}^2 \sum_{s=0}^2 (-1)^{r+s} \binom{2}{r} \binom{2}{s} E_{P_j}[X_k^r Y^s] E_{P_j}[X_k]^{2-r} E_{P_j}[Y]^{2-s} \\ &\quad + \frac{\text{Corr}_{P_j}(X_k, Y)^2}{4} \sum_{r=0}^4 (-1)^r \binom{4}{r} \left[\frac{E_{P_j}[X_k^r] E_{P_j}[X_k]^{4-r}}{s_{P_j}^4(X_k)} + \frac{E_{P_j}[Y^r] E_{P_j}[Y]^{4-r}}{s_{P_j}^4(Y)} \right] \\ &\quad - \frac{\text{Corr}_{P_j}(X_k, Y)}{s_{P_j}(X_k)^3 s_{P_j}(Y)} \sum_{r=0}^3 \sum_{s=0}^1 (-1)^{r+s} \binom{3}{r} E_{P_j}[X_k^r Y^s] E_{P_j}[X_k]^{3-r} E_{P_j}[Y]^{1-s} \\ &\quad - \frac{\text{Corr}_{P_j}(X_k, Y)}{s_{P_j}(X_k) s_{P_j}(Y)^3} \sum_{r=0}^1 \sum_{s=0}^3 (-1)^{r+s} \binom{3}{s} E_{P_j}[X_k^r Y^s] E_{P_j}[X_k]^{1-r} E_{P_j}[Y]^{3-s}. \end{aligned}$$

Observe that all expectations on the right-hand side above are expectations over some $f \in \mathcal{F}_n$ applied to the observed data structure. It follows that the above can be computed in $O(1)$ time using a subset of the $O(p)$ expectation, standard deviation, and correlation estimates stored in H_j . Let **CALCSIGHAT** denote the function which takes as input H_j and d_j and outputs $\hat{\sigma}_{nj}$. We have shown that **CALCSIGHAT**(H_j, d_j) runs in $O(1)$ time.

The pseudocode in **ESTPSI** describes our estimator, with most of the work done in the recursion step described in the function **RECURSION**. Because each call of **RECURSION** runs in $O(p)$ time, the $n - \ell_n = O(n)$ step for loop in **ESTPSI** requires time $O(np)$ time. The storage requirement of each call of **RECURSION** is $O(p)$. Because the code in the for loop in **ESTPSI** deletes the output from the previous recursion step, the total storage requirement of **ESTPSI** is $O(p)$.

Algorithm Recursion Step for Estimating $\Psi(P_0)$

```

function RECURSION( $O_{j+1}, \psi_j, H_j, \bar{\sigma}_j, \ell_n$ )
2:   if  $j < \ell_n$  then  $\psi_{j+1} = 0$  and  $\bar{\sigma}_{j+1} = 0$ 
      else
4:      $d_j = \text{MAXIMIZER}(H_j)$ 
       $\hat{\sigma}_{nj} = \text{CALCSIGHAT}(H_j, d_j)$ 
6:      $D^{d_j}(P_j)(O_{j+1}) = \text{CALCD}(H_j, O_{j+1})$ 
       $\psi_{j+1} = \frac{\psi_j}{\bar{\sigma}_j} + \frac{\text{Corr}_{P_j}(X_{d_j}, Y) + D^{d_j}(P_j)(O_{j+1})}{\hat{\sigma}_{nj}}$  ▷ By convention,  $0/0 = 0$ .
8:      $\bar{\sigma}_{j+1} = \frac{1}{j+1} [j\bar{\sigma}_j + \hat{\sigma}_{nj}]$ 
       $H_{j+1} = \text{UPDATEH}(O_{j+1}, H_j)$ 
10:  return  $(\psi_{j+1}, \bar{\sigma}_{j+1}, H_{j+1})$ 

```

Algorithm Estimate $\Psi(P_0)$ Using Sample of Size n

function ESTPSI(n, ℓ_n)

2: Read O_1, O_2 from data stream

 Base case: $\psi_2 = 0, \bar{\sigma}_2 = 0$, and $H_2 = \text{INITIALIZEH}(O_1, O_2)$

4: **for** $j = 2, \dots, n - 1$ **do**

 Read O_{j+1} from data stream

6: $(\psi_{j+1}, \bar{\sigma}_{j+1}, H_{j+1}) = \text{RECURSION}(O_{j+1}, \psi_j, H_j, \bar{\sigma}_j, \ell_n)$

 Remove $(O_{j+1}, \psi_j, H_j, \bar{\sigma}_j)$ from memory

8: **return** Point estimate ψ_n and confidence interval $[\psi_n \pm 1.96\bar{\sigma}_n/\sqrt{n}]$
